

NEXOF-RA

NESSI Open Framework – Reference Architecture

IST- FP7-216446



**Open Architecture Specification Process
Open Construction Cycle #2**

Security Area

Topic: Multilevel Security for SOA

Position Paper template

Contact: David Llewellyn-Jones, Liverpool John Moores
University

Date of publication: 4th March 2009

Action Required by February 23rd, 2009

To view a copy of this license, visit <http://creativecommons.org/licenses/by/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

1 DETAILS ABOUT THE CONTRIBUTORS

Main Contact / Lead contributor

- Name: Dr David Llewellyn-Jones
- Affiliation: Liverpool John Moores University
- Email: D.Llewellyn-Jones@ljmu.ac.uk
- Phone: +44 (0)151 231 2082
- Mobile Phone: n/a

Additional Contributors (if any)

- Name: Professor Madjid Merabti
- Affiliation: Liverpool John Moores University
- Email: M.Merabti@ljmu.ac.uk
- Phone: +44 (0)151 231 2104
- Mobile Phone: n/a

- Name: Dr Qi Shi
- Affiliation: Liverpool John Moores University
- Email: Q.Shi@ljmu.ac.uk
- Phone: +44 (0)151 231 2272
- Mobile Phone: n/a

- Name: Dr Bob Askwith
- Affiliation: Liverpool John Moores University
- Email: R.J.Askwith@ljmu.ac.uk
- Phone: +44 (0)151 231 2320
- Mobile Phone: n/a

2 TOPIC OF POSITION PAPER

2.1 Title

Managing Risk in Multilevel Security SOAs

2.2 Problem statement/Challenge addressed

Multilevel Secure (MLS) systems are commonly designed to consider security as an order-relation, either via a well-ordering of security levels or a partially ordered lattice. Devices may handle multiple security levels simultaneously as long as the mandatory access control requirements are enforced by the device. The approach taken in the Department of Defense *Orange Book* to determine the risk of compromise (or misuse) of such systems is to assign higher risk to systems able to manipulate higher security levels, as well as higher risk as the range of levels applicable to a device increases [1]. These risks are then translated to the accreditation required by the device. Suggested risk levels have been proposed for individual devices, such as in the Department of Defense *Yellow Book* [2], however the risks associated with device use change in a networked environment where data may be shared between devices. For example, the Department of Defense *Red Book* [3] highlights the cascade vulnerability problem as a particular issue that must be taken into account when multiple devices are networked together.

Such problems have resulted in the need for intricate and cumbersome procedures (such as the Heuristic Algorithm [3]) that – whilst ensuring risks are within acceptable limits – have the effect of hindering the networking of MLS systems. Such procedures are especially problematic where dynamic connections would otherwise be beneficial.

It is reasonable to assume that similar risk and accreditation strategies would be applied to interacting services communicating across the network as components within a Service Oriented Architecture (SOA). The difficulty of calculating risk levels on composed MLS services is likely to cause particular difficulties given the dynamic nature of SOAs and the need to move towards automated service configuration and composition.

The challenge we present is therefore that of designing methods to determine appropriate and acceptable risk levels as services are composed, and how such methods can be applied to services in a SOA.

2.3 Background

The School of Computing and Mathematical Sciences at Liverpool John Moores University has a strong research focus on network security, as well as device and service interoperation generally. The School has been involved in active research in the area of secure component composition for over ten years, with an international track record of conference and journal publications. This work has resulted in a number of nationally funded projects looking into the area, with a more recent focus on secure interoperation in MLS systems.

Liverpool John Moores University has been working in the area of dynamic run-time security analysis as part of a 2-year EPSRC-funded project in collaboration with Thales Research and Technology UK (Ltd). The aim of this project is to develop theory and tools for the analysis of secure interoperability between software components communicating across a network. This combined expertise has borne various results and the development of software for the improvement of security in networked MLS systems.

2.4 Approach proposed (for solution)/Details

The question of how to assess risk in a network, or a MLS network comprised of multiple intercommunicating systems, is not a new one. A variety of approaches have been considered, not least those that try to tackle cascade vulnerability detection and its associated correction problem [4].

In the first instance however, in order to provide security in a MLS system based on the accreditation process, methods are required to determine associated levels of risk.

The Trusted Computer System Evaluation Criteria [1] already specify certain risk correlations for individual systems, for example through the Trusted Computing Base (TCB) rating system. However, with multiple networked devices the interaction between these risks becomes considerably more complex.

A number of different approaches have therefore been proposed for addressing the more complex problem of assigning risk in MLS systems-of-systems. The original approach proposed by the Trusted Network Interpretation of TCSEC involves calculating risk as the maximum of the risks of all devices that would need to be compromised in order for the data downgrade to occur. A more refined analysis has been developed by M. P. Lee [5] that compares risk of data compromise for computers with the equivalent risk as it would ordinarily apply to people. By postulating probability density functions and calculating their convolution through integration Lee establishes a more accurate picture of how risk is affected by devices interacting. Other approaches might also be considered. A more detailed discussion about the various methods for evaluating risk when systems interact is provided in the appendix.

One clear conclusion that we can draw from this is that the risks that could result in illegitimate data downgrades are affected intrinsically by the way services are composed and interact. This then impacts on the security of the system, which may have been accredited to operate only below a certain level of risk with certain types of data. Thus there is a clear need to be able to assess risk when dealing with interacting services.

For an effective approach to considering risk in a MLS SOA, we propose that a number of different automated methods of analysing risk should be available for use. This requirement stems partly from the need for continued research into the most applicable risk analysis approach, for which no categorical answer currently exists, but also from the fact that different contexts are likely to require different risk considerations, or even multiple simultaneous analyses.

Such an approach would entail the integration of risk analysis techniques into the SOA, as a Risk Management Service (RMS) available to be queried by other parts of the architecture as necessary. This would allow other aspects of the architecture, such as those designed to control access based on authorization or prevent unauthorized declassification of material, to operate independently and without the need to calculate risks as an additional consideration.

Our current implementation that makes use of the risk assessment described here has been developed by Thales Research and Technology as an extension to their SUPHICE architecture [6]. Although the existing SUPHICE architecture has been designed for the creation of secure communication links between administrative domains, the IP cryptographic devices used for this are actually exposed as services on the network and could therefore be replaced by other services to reflect a more general SOA. Abstracting the existing architecture and introducing the RMS would result in a configuration similar to that shown in Figure 1. This depicts the process for establishing a connection between two services *A* and *B* (service discovery and other aspects have been omitted for clarity).

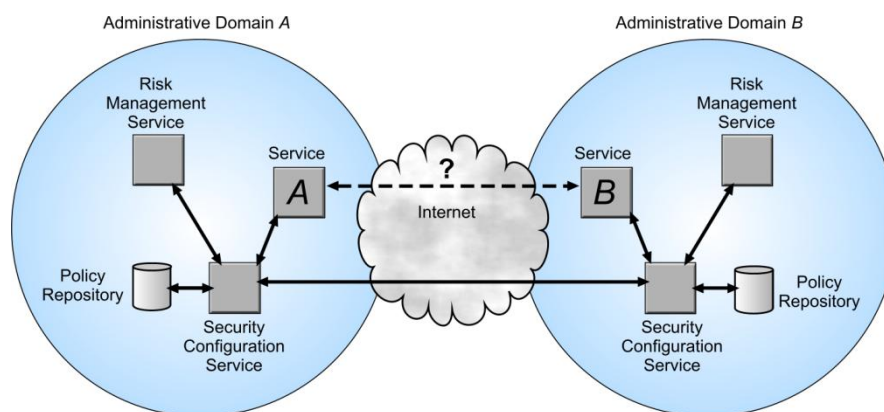


Figure 1. SUPHICE architecture with RMS in SOA context.

The architecture and configuration is ostensibly similar to that of SUPHICE, but with the additional inclusion of the RMS. This is able to perform the risk evaluation based on the services and their connections, and would be queried by the Security Configuration Service in the situation where the risk level appears as a criterion in the domain policy.

The Security Configuration Service acts as a general service for ensuring that other individual services provide adequate security resources on their own. However, in itself this would not tackle the increased risk associated with composition services as outlined here; hence the requirement that the Security Configuration Service liaise with the Risk Management Service when making service configuration decisions.

The Risk Management Service should provide a number of the existing risk assessment methods, with the appropriate technique to be used defined by the

policy for the SOA. In addition, this would also allow new techniques to be incorporated into the Risk Management Service as they are developed.

It should be stressed that although the capability to determine composed risk levels may not represent a functional element within an MLS SOA, without such a capability there remains a very serious danger that the security provisions otherwise provided might become inadequate when applied to multiple interacting services. This is one of the key issues that the cascade vulnerability problem highlights. However, by introducing tools to tackle the problem into the architecture and automatically acting on the resulting risk analysis, configurations that result in unacceptable risk levels can in theory be avoided.

Although a number of algorithms have been proposed for the analysis of risk in networked MLS systems – in particular those attempting to determine or correct the existence of a cascade vulnerability – we know of no other solutions that have been implemented in an automatic manner within MLS networks other than the work on SUPHICE outlined above. This current implementation considers only one variant of the risk analysis, but nonetheless demonstrates that a service as an infrastructure component able to apply a variety of risk evaluations represents an important – yet achievable – goal.

In order to demonstrate the further viability of the approach, additional theoretical risk evaluation methods can be incorporated into the analysis service. Candidate methods might include the Fitch-Hoffman algorithm [4], the Horton Algorithm [4], Simulated Annealing [7], Soft Constraints [8] and so on. Some of these are more suitable for use with a variety of risk calculations (probability density functions) than others. An ideal solution might comprise of a generic framework into which a number of different analysis techniques could be incorporated.

2.5 Dependencies

There are no dependencies identified at this time.

3 SUMMARY

The question of how to determine the risk of compromise of a set of interacting systems is particularly acute in MLS systems, where failure to accurately determine risk can lead to a system with inadequate protection. This applies even more so to MLS Service Oriented Architectures, where multiple services are liable to interact, and where the accreditation of a system is directly dependent on the risks and threats they are subjected to.

A number of strategies have been proposed for determining such risks, from the traditional ‘maximum risk’ approach outlined in the Trusted Network Interpretation of the TCSEC [3], to more refined probabilistic approaches.

In addition, various theoretical algorithms have been developed that allow the risks to be analysed in order to inform an understanding of whether multiple systems can be safely used together.

We therefore propose the application of these algorithms as a Risk Management Service integrated within the architecture to be consulted by the Security Configuration Service. The Risk Management Service should expose functionality to apply the existing algorithms in a flexible manner.

This would allow multiple methods for determining risk to be applied depending on the context. It is a particular characteristic of MLS systems that they require such risk analysis in order to provide adequately security when multiple services interact across the network. Capability to determine this risk is therefore necessary in SOAs in order for other security measures that are applied within the SOA to remain effective.

Although the algorithms that can be applied are generally well understood, there currently are only limited examples of their application, particularly from the perspective of SOAs. The proposed approach would allow the implementation of a solution that is able to develop as our understanding of the risk dynamics of MLS SOAs increases, while at the same time allowing the validity of the approach to be demonstrated through application of existing techniques.

APPENDIX

Even when risks associated with individual systems or services can be known, establishing the consequent risks that would apply to multiple services as they interact requires some careful consideration. We describe a number of possible approaches here.

The traditional approach taken by the Trusted Network Interpretation of TCSEC when discussing the cascade vulnerability problem is to assume that with two (or more) connected devices on a network, the risk of both systems being compromised is equivalent to the risk of the strongest being compromised. The reasoning behind this is that there should never be a situation in which a malicious attacker can “compromise information across a range of security levels that is greater than the accreditation range of any of the component systems he must defeat to do so” [3].

Subsequently a more refined analysis has been performed by M. P. Lee [5] that utilises a comparison with how security clearance is assigned to people. In particular, it considers the notion that the security clearance assigned to an individual is an indication of how likely the *process* is to identify a trustworthy person, rather than how trustworthy the *person* actually is. In addition, the analysis also considers the distinction that can be made between the likelihood of a person with clearance misusing a piece of data and the likelihood of a person with clearance passing it on to someone else improperly.

The semi-rigorous treatment presented suggests that the dominant probability distribution relating to a person causing damage due to holding classified information is as a result of passing the information on improperly, rather than from someone causing damage with that information directly.

Applying this to a computer system, using a suitable probability function can yield a result for the combined probability that two systems might be compromised. This is achieved by hypothesising a suitable probability density function $p_R(t)$ where R is the TCB rating and t is the level of threat. The probability density function $p_J(t)$ for two connected components with the same rating can then be calculated as the convolution integral

$$p_J(t) = \int_{-\infty}^{\infty} p_R(x)p_R(t-x) dx.$$

Lee takes a probability density function of $p_R(t) = ae^{-bt}$ as in Figure 2(a) to yield a combined probability of $p_J(t) = a^2te^{-bt}$ as in Figure 2(b).




Figure 2. (a) $p_R(t) = ae^{-bt}$. (b) $p_J(t) = a^2 te^{-bt}$.

This conclusion is different to that of the traditional method. For example, it produces the result that the resistance to threat of two B2 systems is approximately the same as or better than that of a single B3 system, rather than remaining at the B2 level that would be the maximum of the two.

While the results provided by Lee are not conclusive, they do support the conclusion that some service compositions that might ordinarily be classified as ‘too risky’ should actually be considered as having acceptable levels of risk.

Further assessments of overall risk may also be relevant in the context of MLS services. In a networked environment the risk of compromise can be dependent on the level of exposure. From a Service Oriented Architecture perspective we can consider this as a function of the number of services being used to process a particular piece of data. For a set of linearly networked services s_1, \dots, s_n for example, with probability of compromise p_1, \dots, p_n , the overall probability that the system will be compromised can be taken as

$$1 - \prod_{1 \leq i \leq n} (1 - p_i).$$

It should be noted that this effectively constitutes a simplified form of the previously described convolution integral, but using an alternative (simpler) probability that isn’t necessarily taken to be a function of threat level, but rather as a property that can be assigned to each individual service. Nonetheless this may be an appropriate and realistic way to consider threat levels in a Service Oriented Architecture. Again, here we see an alternative (but not necessarily conflicting) way to consider the risk associated with a collection of networked MLS services.

REFERENCES

- [1] Department of Defense, "Trusted Computer System Evaluation Criteria (TCSEC) ("Orange Book")," National Computer Security Center, Ft. Meade, MD 20755, DoD 5200.28-STD, 26 December 1985.
- [2] Department of Defense, "Technical Rational Behind CSC-STD-003-85: Computer Security Requirements -- Guidance for Applying the Department of Defense TCSEC in Specific Environments ("Yellow Book")," National Computer Security Center, Ft. Meade, MD 20755, CSC-STD-004-85, 25 June 1985.
- [3] Department of Defense, "Trusted Network Interpretation of the TCSEC (TNI) ("Red Book")," National Computer Security Center, Ft. Meade, MD 20755, NCSC-TG-005, 31 July 1987.
- [4] J. D. Horton, R. H. Cooper, W. F. Hyslop, B. G. Nickerson, O. K. Ward, R. Harland, E. Ashby, and W. M. Stewart, "The cascade vulnerability problem," *Journal of Computer Security*, vol. 2(4), pp. 279-90, 24-26 May 1993.
- [5] T. M. P. Lee, "Statistical models of trust: TCBs vs. people," in Proceedings 1989 IEEE Symposium on Security and Privacy (Cat. No.89CH2703-7), Oakland, CA, USA, 1-3 May 1989.
- [6] S. Butler and M. Irons, "Secure Unplanned Provisioning of High Integrity Communications across Europe - D3.1 Architecture Definition," D31-O-PU-003-TRT, 12 June 2006.
- [7] S. Gritzalis and D. Spinellis, "The cascade vulnerability problem: the detection problem and a simulated annealing approach for its correction," *Microprocessors and Microsystems*, vol. 21(10), pp. 621-7, 30 April 1998.
- [8] S. Bistarelli, S. N. Foley, and B. O. Sullivan, "A soft constraint-based approach to the cascade vulnerability problem," *Journal of Computer Security*, vol. 13(5), pp. 699-720, 2005.